

Bridging the Gap: a Semantic Similarity Measure between Queries and Documents

Sun Kim, W. John Wilbur and Zhiyong Lu

National Center for Biotechnology Information
National Library of Medicine, National Institutes of Health
Bethesda, MD 20894, USA
{sun.kim, john.wilbur, zhiyong.lu}@nih.gov

Abstract

The main approach of traditional information retrieval (IR) is to examine how many words from a query appear in a document. A drawback of this approach, however, is that it may fail to detect relevant documents where no or only few words from a query are found. The semantic analysis methods such as LSA (latent semantic analysis) and LDA (latent Dirichlet allocation) have been proposed to address the issue, but their performance is not superior compared to common IR approaches. Here we present a query-document similarity measure motivated by the Word Mover's Distance. Unlike other similarity measures, the proposed method relies on neural word embeddings to calculate the distance between words. Our method is efficient and straightforward to implement. The experimental results on TREC and PubMed® show that our approach provides significantly better performance than BM25. We also discuss the pros and cons of our approach and show that there is a synergy effect when the word embedding measure is combined with the BM25 function.

1 Introduction

In information retrieval (IR), queries and documents are typically represented by term vectors where each term is a content word and weighted by *tf-idf* or other weighting schemes (Salton and Buckley, 1988). The similarity of a query and a document is then determined as a dot product or cosine similarity. Although this works reasonably, the traditional IR scheme often fails to find relevant documents when synonymous or polysemous words are used in a dataset, e.g. a document including only “neoplasm” cannot be found

when the word “cancer” is used in a query. One solution of this problem is to use query expansion (Carpineto and Romano, 2012) or dictionaries, but these alternatives still depend on the same philosophy, i.e. queries and documents should share exactly the same words.

While the term vector model computes similarities in a sparse and high-dimensional space, the semantic analysis methods such as latent semantic analysis (LSA) (Deerwester et al., 1990; Hofmann, 1999) and latent Dirichlet allocation (LDA) (Blei et al., 2003) learn dense vector representations in a low-dimensional space. These methods choose a vector embedding for each term and estimate a similarity between terms by taking an inner product of their corresponding embeddings (Sordani et al., 2014). Since the similarity is calculated in a latent (semantic) space based on context, the semantic analysis approaches do not require having common words between a query and documents. However, it has been shown that LSA and LDA methods do not produce superior results in various IR tasks (Maas et al., 2011; Baroni et al., 2014; Pennington et al., 2014) and the classic ranking method, BM25 (Robertson and Zaragoza, 2009), even outperforms them in document ranking (Atreya and Elkan, 2011; Nalisnick et al., 2016).

Neural word embedding (Bengio et al., 2003; Mikolov et al., 2013) is similar to the semantic analysis methods described above. The main difference is that LSA and LDA utilize co-occurrences of words but neural word embedding learns to predict context. Moreover, training of semantic vectors is derived from neural networks. The neural word embedding models have gained popularity in recent years due to their high performance

in NLP tasks (Levy and Goldberg, 2014). Here we present a query-document similarity measure using a neural word embedding approach. This work is particularly motivated by the Word Mover’s Distance (Kusner et al., 2015). Unlike the common similarity measure taking query/document centroids of word embeddings, the proposed method evaluates a distance between individual words from a query and a document. The experimental results on TREC (Hersh et al., 2006; Hersh et al., 2007) and PubMed¹ show that our approach is significantly better than the BM25 ranking. We also show that an additional improvement can be achieved by combining the word embedding approach and the BM25 function.

Taken together, we make the following major contributions in this work. First, to our best knowledge, this work represents the first investigation of query-document similarity for information retrieval using the recently proposed Word Mover’s Distance. Second, we modified the original Word Mover’s Distance algorithm so that it is computationally less expensive and thus more practical for real-world search scenarios (e.g. biomedical literature search). Third, we propose a novel strategy for using PubMed article titles as a proxy for relevance so that ranking algorithms can be evaluated and compared at a large scale without the involvement of tedious manual relevance judgment process. Finally, on both TREC and PubMed datasets, our proposed method achieved stronger performance compared to BM25.

2 Methods

A common approach of computing a similarity between texts (e.g. phrases, sentences or documents) is to take a centroid of word embeddings, and evaluate an inner product or cosine similarity between centroids² (Furnas et al., 1988; Nalisnick et al., 2016). This has found use in classification and clustering because they require to find an overall topic of each document. However, taking a simple centroid is not a good approximator for calculating a distance between a document and a query (Kusner et al., 2015).

¹<https://pubmed.gov>

²The implementation of *word2vec* also uses centroids of word vectors for calculating similarities (<https://code.google.com/archive/p/word2vec>).

A standard method of document search, on the other hand, is to find a set of documents that include query words or a subtopic relevant to the query. Consistent with this, our approach here is to measure the distance between individual words, not the average distance between a query and a document.

2.1 Word Mover’s Distance

Our work is based on the Word Mover’s Distance between text documents (Kusner et al., 2015), which calculates the minimum cumulative distance that words from a document need to travel to match words from a second document. First, we might assume that documents are represented by normalized bag-of-words (BOW) vectors, i.e. if a word w_i appears tf_i times in a document, the weight $d_i = \frac{tf_i}{\sum_{j=1}^n tf_j}$, where n is number of words in the document. The higher the weight, the more important the word. Combined with a measure of word importance, we also require a method to measure the relatedness of a pair of words. For this purpose, we use a neural embedding approach. The dissimilarity c between w_i and w_j can be calculated by $c(i, j) = \|\mathbf{x}_i - \mathbf{x}_j\|_2$, where \mathbf{x}_i and \mathbf{x}_j are the embeddings of w_i and w_j , respectively. The Word Mover’s Distance makes use of word importance and the relatedness of words as we now describe.

Let \mathbf{D} and \mathbf{D}' be BOW representations of two documents D and D' . Let $\mathbf{T} \in \mathcal{R}^{n \times n}$ be a flow matrix, where $\mathbf{T}_{ij} \geq 0$ denotes how much it costs to travel from w_i in D to w_j in D' , and n is the number of unique words appearing in D and D' . To entirely transform \mathbf{D} to \mathbf{D}' , we ensure that the entire outgoing flow from w_i equals d_i and the incoming flow to w_j equals d'_j . The Word Mover’s Distance between D and D' is then defined as the minimum cumulative cost required to move all words from D to D' , i.e.

$$\begin{aligned} \min_{\mathbf{T} \geq 0} \quad & \sum_{i,j=1}^n \mathbf{T}_{ij} c(i, j) \\ \text{subject to} \quad & \sum_{j=1}^n \mathbf{T}_{ij} = d_i, \forall i \in \{1, \dots, n\} \\ & \sum_{i=1}^n \mathbf{T}_{ij} = d'_j, \forall j \in \{1, \dots, n\}. \end{aligned} \tag{1}$$

The solution is attained by finding \mathbf{T}_{ij} that minimizes the expression in (1). Kusner et al. (2015) applied this to obtain nearest neighbors for document classification, i.e. k -NN classification and it produced outstanding performance among other state-of-the-art approaches. What we have just described is the approach given in Kusner et al. We will modify the word weights and the measure of the relatedness of words to better suit our application.

2.2 Our Query-Document Similarity Measure

While the prior work gives a hint that the Word Mover’s Distance is a reasonable choice for evaluating a similarity between documents, it is uncertain how the same measure could be used for searching documents from a query. First, it is expensive to compute the Word Mover’s Distance. The time complexity of solving the distance problem is $O(n^3 \log n)$ (Pele and Werman, 2009). Second, the semantic space of queries is not the same as those of documents. A query consists of a small number of words in general, hence words in a query tend to be more abstract and somewhat vague. On the contrary, a text document is longer and more informational. Having this in mind, we realize that two distinctive components are required for query-document search: 1) mapping queries to documents using a word embedding model trained on a document set and 2) mapping documents to queries using a word embedding model obtained from a query set. In this work, however, we aim to address the former, and the mapping of documents to queries remains as future work.

For our purpose, we will change the word weight d_i to incorporate inverse document frequency (*idf*), i.e. $d_i = idf(i) \frac{tf_i}{\sum_{j=1}^n tf_j}$, where $idf(i) = \log \frac{K - k_i + 0.5}{k_i + 0.5}$. K is the size of a document set and k_i is the number of documents that include the i th term. This is the *idf* factor normally used in *tf-idf* and BM25 (Witten et al., 1999; Wilbur, 2001).

Let \mathbf{Q} and \mathbf{D} be BOW representations of a query Q and a document D . D and D' in Section 2.1 are now replaced by Q and D , respectively. The word embedding model is trained on a set of (large) documents. Since we want to have a higher score for documents relevant to Q , $c(i, j)$ is redefined as a cosine similarity, i.e. $c(i, j) = \frac{\mathbf{x}_i \cdot \mathbf{x}_j}{\|\mathbf{x}_i\| \|\mathbf{x}_j\|}$. In addition,

Dataset	# Documents
TREC	162,259
PubMed	15,791,442

Table 1: Number of documents in the TREC and PubMed datasets.

the problem we try to solve is the flow $Q \rightarrow D$. Hence, Eq. 1 is rewritten as follows.

$$\begin{aligned} \max_{\mathbf{T} \geq 0} \quad & \sum_{i,j=1}^n \mathbf{T}_{ij} c(i, j) \\ \text{subject to} \quad & \sum_{j=1}^n \mathbf{T}_{ij} = d_i, \forall i \in \{1, \dots, n\}. \end{aligned} \quad (2)$$

The optimal solution of the expression in (2) is to map each word in Q to the most similar word in D based on word embeddings. The time complexity for getting the optimal solution is $O(mn)$, where m is the number of unique query words and n is the number of unique document words. In general, $m \ll n$ and evaluating the similarity between a query and documents can be implemented in parallel computation. Thus, the document ranking process can be quite efficient.

3 Experimental Results

To evaluate our word embedding approach, we used two scientific literature datasets: TREC and PubMed. TREC is the benchmark set created for TREC 2006 and 2007 Genomics Tracks (Hersh et al., 2006; Hersh et al., 2007). The original task is to retrieve passages relevant to topics (i.e. queries) from full-text articles, but the same set can be utilized for searching relevant documents. Our setup is more challenging because we only use PubMed abstracts, not full-text articles, to find evidence. The PubMed set is larger, and it contains all the documents that have non-empty titles and abstracts in PubMed. Table 1 shows the number of documents in each dataset. While the TREC and PubMed sets share essentially the same type of documents, tested queries are quite different. The queries in TREC are a question type, e.g. “what is the role of MMS2 in cancer?” However, the PubMed set uses actual queries from PubMed users.

We used the skip-gram model of *word2vec* (Mikolov et al., 2013) to obtain word embeddings.

Dataset	# Topics	BM25	EMBED
TREC 2006	26	0.3136	0.3732
TREC 2007	36	0.2463	0.2601

Table 2: Average precision of BM25 and our embedding approach (EMBED) on the TREC set.

word2vec was trained on titles and abstracts from over 25 million PubMed documents³, and the trained model is available online⁴.

For experiments, we removed stopwords from queries and documents. BM25⁵ is chosen for performance comparison in following subsections. Among document ranking functions, BM25 shows a competitive performance (Trotman et al., 2014). It also outperforms co-occurrence based word embedding models (Atreya and Elkan, 2011; Nalisnick et al., 2016).

3.1 TREC Dataset

Table 2 presents the average precision of BM25 and our embedding approach on the TREC dataset. Relevance judgements in TREC are based on the pooling method (Manning et al., 2008), i.e. relevance is manually assessed for top ranking documents returned by participating systems. Therefore, we only used the documents that annotators reviewed for evaluation (Lu et al., 2009). As shown in the table, the embedding approach boosts the average precision of BM25 by 19% and 6% on TREC 2006 and 2007, respectively.

3.2 PubMed Dataset

Even if TREC is a reasonable set to evaluate search systems, queries used in the set do not reflect an actual search scenario in PubMed. For manual judgements, TREC queries should be specific and unambiguous if possible. This constraint makes queries longer than usual PubMed queries. The average lengths⁶ of TREC 2006 and 2007 queries are 6.8 and 5.7 words, respectively. For this reason, we examined the search performance over 19119 PubMed

³*word vector size* = 100 and *window size* = 10. The parameters were optimized to produce high recall for synonyms. Note that an independent set was used for tuning *word2vec* parameters.

⁴<https://www.ncbi.nlm.nih.gov/IRET/DATASET>

⁵ $k = 1.9$ and $b = 1.0$ (Lin and Wilbur, 2007)

⁶Average number of words after removing stopwords.

Dataset	# Queries	BM25	EMBED
Query words = 2	10485	0.1823	0.2242
Query words = 3	6151	0.1292	0.2056
Query words > 3	2483	0.0921	0.1986

Table 3: Average precision of BM25 and our embedding approach (EMBED) on the PubMed set.

Dataset	# Topics	BM25	EMBED
TREC 2006 (all)	26	0.2140	0.2510
TREC 2007 (all)	36	0.1447	0.1253

Table 4: Average precision of BM25 and our approach (EMBED) on the TREC set. All documents in the set are scored.

user queries⁷. Manual judgements are not available for PubMed queries, hence we used the title of each PubMed document⁸ as a proxy for relevance, i.e. a document is positive if its title includes all query words. The search is performed only on abstracts. This evaluation is based on the observation that humans can recognize the subject of a document satisfactorily from the title (Resnick, 1961).

Table 3 shows the average precision of BM25 and the embedding approach over user queries of different lengths. When using the embedding approach, the average precision is increased by 23%, 59% and 116% for queries with 2, 3, and 4 or more words, respectively. Interestingly, the BM25 performance degrades for longer queries, but the performance does not decrease much in the word embedding approach. We speculate that BM25 picks up more irrelevant abstracts because it scores based only on the exact words found, whereas the word embedding approach produces a semantic score even for query words that do not appear in the document.

3.3 Word Embeddings with BM25

Another experiment we performed was to score all documents from TREC using BM25 and the word embedding approach. In this experiment, only positives are used from manual judged labels, and we assume all the remaining documents are negatives (Webber and Park, 2009). As expected, the performance drops significantly (Table 4) compared to the results in Section 3.1. For TREC

⁷These are frequent user queries entered between Jun. 1 and Jun. 30, 2015.

⁸PubMed documents are semi-structured, and each document has a title and an abstract.




Figure 1: TREC performance changes depending on α .

2007, the word embedding approach performs worse than BM25. This may be due to the bias issue of the pooling method (Buckley et al., 2007; Webber and Park, 2009) or simply added noises by scoring all documents. Since we found a case that BM25 performed better, we considered a linear combination of the word embedding and BM25, i.e. $\alpha \cdot \text{EMBED} + (1 - \alpha) \cdot \text{BM25}$. As depicted in Figure 1, any α between 0.1 and 0.9 yields better performance than BM25 only. This demonstrates that there is a synergy effect by merging the two approaches.

4 Conclusions

We presented a word embedding approach for measuring a similarity between a query and a document. Starting from the Word Mover’s Distance, we reinterpreted the model for a query-document search problem. Even with the $Q \rightarrow D$ flow only, the word embedding approach is already efficient and effective. In this setup, the proposed approach cannot distinguish documents when they include all query words, but surprisingly, the word embedding approach shows remarkable performance on TREC and PubMed datasets. Our immediate plan is to analyze the $D \rightarrow Q$ flow by learning word embeddings from billions of PubMed user queries.

Acknowledgments

This research was supported by the Intramural Research Program of the NIH, National Library of Medicine.

References

- [Atreya and Elkan2011] A. Atreya and C. Elkan. 2011. Latent semantic indexing (LSI) fails for TREC collections. *ACM SIGKDD Explorations Newsletter*, 12(2):5–10.
- [Baroni et al.2014] M. Baroni, G. Dinu, and G. Kruszewski. 2014. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proc. Annual Meeting of the Association for Computational Linguistics (ACL 2014)*, pages 238–247, Jun.
- [Bengio et al.2003] Y. Bengio, R. Ducharme, P. Vincent, and C. Janvin. 2003. A neural probabilistic language model. *Journal of Machine Learning Research*, 3:1137–1155.
- [Blei et al.2003] D. M. Blei, A. Y. Ng, and M. I. Jordan. 2003. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- [Buckley et al.2007] C. Buckley, D. Dimmick, I. Soboroff, and E. Voorhees. 2007. Bias and the limits of pooling for large collections. *Information Retrieval*, 10(6):491–508.
- [Carpineto and Romano2012] C. Carpineto and G. Romano. 2012. A survey of automatic query expansion in information retrieval. *ACM Computing Surveys*, 44(1):1:1–1:50.
- [Deerwester et al.1990] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407.
- [Furnas et al.1988] G. W. Furnas, S. Deerwester, S. T. Dumais, T. K. Landauer, R. A. Harshman, L. A. Streeter, and K. E. Lochbaum. 1988. Information retrieval using a singular value decomposition model of latent semantic structure. In *Proc. Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 465–480, Jun.
- [Hersh et al.2006] W. Hersh, A. M. Cohen, P. Roberts, and H. K. Rekapalli. 2006. TREC 2006 Genomics Track overview. In *Proc. Text REtrieval Conference 2006*, Nov.
- [Hersh et al.2007] W. Hersh, A. M. Cohen, L. Ruslen, and P. Roberts. 2007. TREC 2007 Genomics Track overview. In *Proc. Text REtrieval Conference 2007*, Nov.
- [Hofmann1999] T. Hofmann. 1999. Probabilistic latent semantic indexing. In *Proc. Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 50–57.
- [Kusner et al.2015] M. J. Kusner, Y. Sun, N. I. Kolkin, and K. Q. Weinberger. 2015. From word embeddings

- to document distances. In *Proc. International Conference on Machine Learning (ICML 2015)*, volume 37, pages 957–966, Jul.
- [Levy and Goldberg2014] O. Levy and Y. Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in Neural Information Processing Systems (NIPS 2014)*, pages 2177–2185, Dec.
- [Lin and Wilbur2007] J. Lin and W. J. Wilbur. 2007. PubMed related articles: a probabilistic topic-based model for content similarity. *BMC Bioinformatics*, 8:423.
- [Lu et al.2009] Z. Lu, W. Kim, and W. J. Wilbur. 2009. Evaluation of query expansion using MeSH in PubMed. *Information Retrieval*, 12(1):69–80.
- [Maas et al.2011] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. 2011. Learning word vectors for sentiment analysis. In *Proc. Annual Meeting of the Association for Computational Linguistics (ACL 2011)*, pages 142–150.
- [Manning et al.2008] C. D. Manning, P. Raghavan, and H. Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA.
- [Mikolov et al.2013] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems (NIPS 2013)*, pages 3111–3119, Dec.
- [Nalisnick et al.2016] E. Nalisnick, B. Mitra, N. Craswell, and R. Caruana. 2016. Improving document ranking with dual word embeddings. In *Proc. International World Wide Web Conference (WWW 2016)*, pages 83–84, Apr.
- [Pele and Werman2009] O. Pele and M. Werman. 2009. Fast and robust earth mover’s distances. In *Proc. IEEE International Conference on Computer Vision (ICCV 2009)*, pages 460–467, Sep.
- [Pennington et al.2014] J. Pennington, R. Socher, and C. Manning. 2014. Glove: Global vectors for word representation. In *Proc. Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, pages 1532–1543, Oct.
- [Resnick1961] A. Resnick. 1961. Relative effectiveness of document titles and abstracts for determining relevance of documents. *Science*, 134(3484):1004–1006.
- [Robertson and Zaragoza2009] S. Robertson and H. Zaragoza. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- [Salton and Buckley1988] G. Salton and C. Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information Processing & Management*, 24(5):513–523.
- [Sordoni et al.2014] A. Sordoni, Y. Bengio, and J.-Y. Nie. 2014. Learning concept embeddings for query expansion by quantum entropy minimization. In *Proc. AAAI Conference on Artificial Intelligence*, pages 1586–1592.
- [Trotman et al.2014] A. Trotman, A. Puurula, and B. Burgess. 2014. Improvements to BM25 and language models examined. In *Proc. 2014 Australasian Document Computing Symposium (ADCS 2014)*, pages 58:58–58:65, Nov.
- [Webber and Park2009] W. Webber and L. A. F. Park. 2009. Score adjustment for correction of pooling bias. In *Proc. International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 444–451, Jul.
- [Wilbur2001] W. J. Wilbur. 2001. Global term weights for document retrieval learned from trec data. *Journal of Information Science*, 27(5):303–310.
- [Witten et al.1999] I. H. Witten, A. Moffat, and T. C. Bell. 1999. *Managing Gigabytes*. Morgan-Kaufmann, San Francisco, CA, USA.